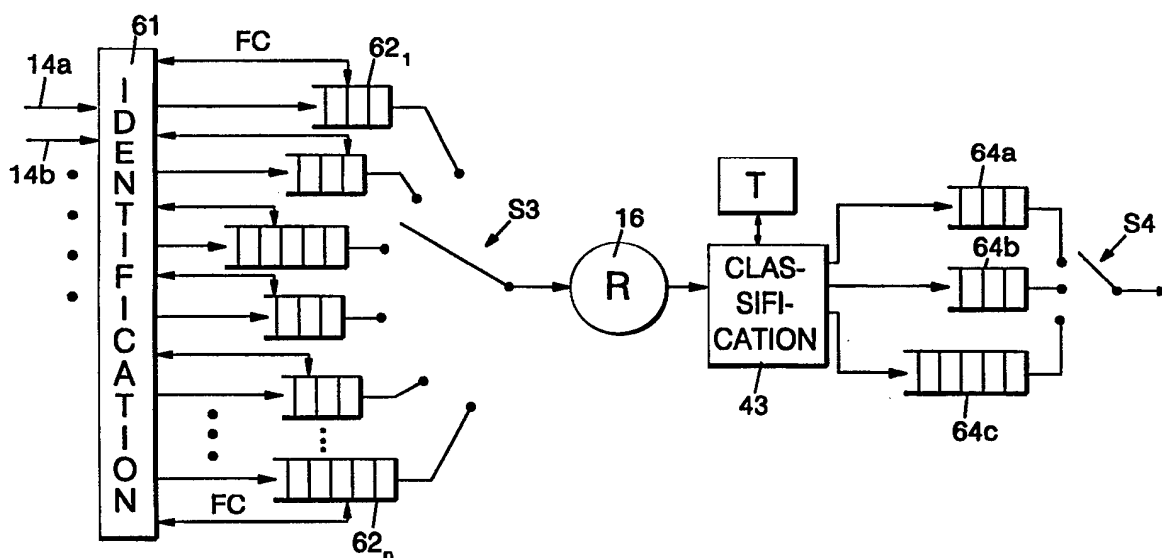




INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁶ : H04L 12/56	A1	(11) International Publication Number: WO 95/15636 (43) International Publication Date: 8 June 1995 (08.06.95)
(21) International Application Number: PCT/FI94/00534 (22) International Filing Date: 29 November 1994 (29.11.94) (30) Priority Data: 935364 30 November 1993 (30.11.93) FI (71) Applicant (for all designated States except US): NOKIA TELECOMMUNICATIONS OY [FI/FI]; Mäkkylän puistotie 1, FIN-02600 Espoo (FI). (72) Inventors; and (75) Inventors/Applicants (for US only): PAJUVIRTA, Juha [FI/FI]; Pihlajatie 29 A 8, FIN-00270 Helsinki (FI). SALOVUORI, Heikki [FI/FI]; Fregattikuja 3, FIN-00850 Helsinki (FI). MATKASELKÄ, Jorma [FI/FI]; Nuijatie 4 C 24, FIN-01650 Vantaa (FI). KASSLIN, Mika [FI/FI]; Torkkelinkatu 17 A 18, FIN-00500 Helsinki (FI). PYHÄLAMMI, Seppo [FI/FI]; Sankaritie 4 A 1, FIN-00320 Helsinki (FI). OLKKONEN, Mikko [FI/FI]; Muurarinkuja 1 D 41, FIN-02600 Espoo (FI). FEHLMAN, Richard [CH/FI]; Kauppalaantie 27-29 A 9, FIN-00320 Helsinki (FI). LAIHO, Mikko [FI/FI]; Ruokoushuoneentie 4, FIN-00730 Helsinki (FI).		(74) Agent: OY KOLSTER AB; Iso Roobertinkatu 23, P.O. Box 148, FIN-00121 Helsinki (FI). (81) Designated States: AM, AT, AU, BB, BG, BR, BY, CA, CH, CN, CZ, DE, DK, EE, ES, FI, GB, GE, HU, JP, KE, KG, KP, KR, KZ, LK, LR, LT, LU, LV, MD, MG, MN, MW, NL, NO, NZ, PL, PT, RO, RU, SD, SE, SI, SK, TJ, TT, UA, US, UZ, VN, European patent (AT, BE, CH, DE, DK, ES, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, ML, MR, NE, SN, TD, TG), ARIPO patent (KE, MW, SD, SZ). Published <i>With international search report.</i> <i>Before the expiration of the time limit for amending the claims and to be republished in the event of the receipt of amendments.</i> <i>In English translation (filed in Finnish).</i>

(54) Title: CONTROL OF OVERLOAD SITUATIONS IN FRAME RELAY NETWORK



(57) Abstract

The invention relates to a method for congestion management in an FR network. The method comprises determining the virtual channel associated with a frame to be transmitted when it is received from a subscriber connection at a subscriber node. In order for the load in the network to be reduced, (a) data is buffered at the input boundary of the subscriber node to virtual-channel-specific buffers (62₁...62_n), (b) the fill rate of said buffers is monitored continuously, and (c) when a frame is received in a virtual-channel-specific buffer which is full, substantially all the contents of the buffer are discarded.

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AT	Austria	GB	United Kingdom	MR	Mauritania
AU	Australia	GE	Georgia	MW	Malawi
BB	Barbados	GN	Guinea	NE	Niger
BE	Belgium	GR	Greece	NL	Netherlands
BF	Burkina Faso	HU	Hungary	NO	Norway
BG	Bulgaria	IE	Ireland	NZ	New Zealand
BJ	Benin	IT	Italy	PL	Poland
BR	Brazil	JP	Japan	PT	Portugal
BY	Belarus	KE	Kenya	RO	Romania
CA	Canada	KG	Kyrgyzstan	RU	Russian Federation
CF	Central African Republic	KP	Democratic People's Republic of Korea	SD	Sudan
CG	Congo	KR	Republic of Korea	SE	Sweden
CH	Switzerland	KZ	Kazakhstan	SI	Slovenia
CI	Côte d'Ivoire	LI	Liechtenstein	SK	Slovakia
CM	Cameroon	LK	Sri Lanka	SN	Senegal
CN	China	LU	Luxembourg	TD	Chad
CS	Czechoslovakia	LV	Latvia	TG	Togo
CZ	Czech Republic	MC	Monaco	TJ	Tajikistan
DE	Germany	MD	Republic of Moldova	TT	Trinidad and Tobago
DK	Denmark	MG	Madagascar	UA	Ukraine
ES	Spain	ML	Mali	US	United States of America
FI	Finland	MN	Mongolia	UZ	Uzbekistan
FR	France			VN	Viet Nam
GA	Gabon				

Control of overload situations in frame relay network

The invention relates to a method according to the preamble of appended claim 1 for congestion management in a frame relay network, and to a subscriber node according to the preamble of appended claim 3.

Congestion means a situation in which the number of transmission requests exceeds the transmission capacity at a certain network point (called a bottle-neck resource) at a specific time. Congestion usually results in overload conditions, as a result of which the buffers overflow, for instance, and so packets will be retransmitted either by the network or the subscriber. The function of congestion management (CM) is to maintain a balance between the transmission requests and the transmission capacity so that the bottle-neck resources operate on an optimal level, and the subscribers are offered service in a way that assures fairness.

Congestion management can be divided into congestion avoidance (CA) and congestion recovery (CR). Congestion avoidance methods aim at preventing the occurrence of congestion in the network by dynamically adjusting the bandwidth of the subscribers in accordance with network congestion conditions and/or by altering the network routing so that part of the traffic load of the bottle-neck resources is shifted to idle resources. The purpose of recovery methods, in turn, is to restore the operation of the bottle-neck resources to an optimal level if the avoidance methods have failed to prevent the occurrence of congestion.

The frame relay (FR) technique is a packet-switched network technique used for the transmission of frames of varying length in place of the packet-switched network connections presently in use. The protocol

(X.25) applied generally in the present packet-switched networks requires plenty of processing and the transmission equipment is expensive, as a result of which the speeds also remain low. These matters are due to the fact that the X.25 standard was developed when the transmission connections used were still rather prone to transmission errors. The starting point of the frame relay technique was a considerably lower transmission line error probability. It has therefore been possible to discard a number of unnecessary functions in the frame relay technique, which makes the delivery of frames rapid and efficient. The Frame Mode Bearer Service is described generally in CCITT recommendation I.233 (Reference 1) and the associated protocol in recommendation Q.922 (Reference 2). Congestion in an FR network and congestion management mechanisms are described in CCITT recommendation I.370 (Reference 3). For a more detailed description of the FR technique, reference is made to *An Overview of Frame Relay Technology*, Datapro Management of Data Communications, McGraw-Hill Incorporated, April 1991, (Reference 4) as well as the above-mentioned recommendations.

As a last resort in cases of extremely severe congestion in a frame relay network, frames can be discarded. Situations of this kind arise when both the transmission capacity of the network and the buffering capacity of a single node are exceeded. In the congestion management mechanisms presently used, frames are discarded almost randomly both on the periphery of the network (at subscriber nodes) and in the middle of the network (at trunk nodes). A frame is discarded if it cannot be buffered due to lack of space, or if the fill rate of the buffer is sufficiently high and the frame contains an indication which allows it to be discarded in the event of congestion, for instance.

Nothing is done to affect the buffer itself; it continues to be emptied to the data link as normal. If the frame is not discarded until at a congested node, the other resources of the network are loaded in vain.

5 The frame must be switched through the network up to the point where it is discarded, i.e. to the congested node, which means wasting the resources.

The object of the present invention is to obviate the drawback described above and to provide a new type of congestion management method for a frame relay network, said method allowing unnecessary loading of the network to be avoided. This is achieved with a method according to the invention, which is characterized by what is disclosed in the characterizing part of

10 appended claim 1. The FR network subscriber node according to the invention is characterized by what is disclosed in the characterizing part of appended claim

15 3.

The idea of the invention is to empty the contents of an entire virtual-channel-specific buffer at a time on the periphery of the network at a subscriber node where the frames are buffered to virtual-channel-specific buffers. (The virtual channel refers to a virtual connection portion having the length

20 of one transmission link while the virtual connection is the actual packet-switched end-to-end FR connection.)

25

The method according to the invention makes it possible to avoid discarding random frames, on account of which the applications do not have to retransmit packet sequences as often as before.

30

In the following, the invention and its preferred embodiments will be described in greater detail with reference to the examples illustrated in the accompanying drawings, in which

Figure 1 illustrates a typical operating environment of the method according to the invention,

Figure 2 illustrates an FR network subscriber node according to the invention,

5 Figure 3 illustrates the format of a frame to be delivered in an FR network, and

Figure 4 illustrates a trunk node of an FR network employing the method according to the invention.

10 A frame relay network can be used by several different applications, which do not require similar services. It is therefore advantageous to employ the method of the invention in a network where the services are divided into different classes according to the applications, taking into account the two most significant parameters (frame loss probability and delay). Such
15 a solution is disclosed in Finnish Patent Application No. 925671. In this application, it is suggested that the services be divided into three classes as follows:

- the first service class (class 1) offers
20 interactive service where the delay is kept short,
- the second service class (class 2) offers a low frame loss probability without the delay having any great significance, and
- the third service class (class 3) offers both
25 a short delay and a low frame loss probability.

Each trunk connection of a network realized in this way has service-class-specific buffers, one for each service class. A subscriber node, in turn, has virtual-channel-specific buffers on the side of the
30 subscriber interface. These solutions will be described more closely below; otherwise reference is made to the Finnish patent application cited above.

(Subscriber) applications using a frame relay network send data towards the network in frames, typically
35 a so-called window at a time. If the application

and the protocol used by it detect that a frame included in a window has been lost, the entire window is usually retransmitted instead of transmitting only the frame that has been lost. Thus the loss of one frame in the network affects the application in the same way as the loss of an entire window. In each case, the application protocol must retransmit a whole windowful of frames.

Figure 1 shows an FR network offering public network services, that is, a frame relay network 12 interconnecting local area networks 11 of different offices A...E of a single corporation or a plurality of corporations. The local area network 11 of each office has access to the FR service via a local area network bridge 13 and a data link indicated with the references 14a...14e, respectively. The connection between an FR subscriber A...E and an FR network node N is known *per se*, wherefore it will not be described more closely herein. More detailed information about local area networks and bridges used in their interconnection can be found e.g. in an article by Michael Grimshaw *LAN Interconnections Technology*, Telecommunications, February 1991, and in *Lähiverkko-opas*, Leena Jaakonmäki, Suomen ATK-kustannus Oy, 1991, which are incorporated herein by reference.

A typical feature of the known node structure of the FR network is that the same buffer is used for all frames, assuming that they are routed to the same physical connection. According to the present invention, on the contrary, buffers corresponding to the above-described service classes are preferably provided at the output boundary of all network nodes and at the input boundary having trunk connections. Figure 2 illustrates this kind of solution at a trunk node in the network. The node receives FR frames originally assembled in the bridges 13 of the subscriber connections (Figure 1). The

frame of the subscriber LAN 11 is inserted into the information field of the FR frame in the bridge 13 (with the exception of timing bits and other similar bits). Figure 3 illustrates the insertion of a LAN frame 38 into the information field of an FR frame 39. It also shows a typical FR network frame format where the address field preceding the information field comprises two octets (bits 1 to 8). The bits 3 to 8 of the first octet and the bits 5 to 8 of the second octet form a data link connection identifier DLCI, which indicates to the node e.g. the virtual connection and virtual channel to which a particular frame belongs. The virtual channels are distinguished from each other by means of the data link connection identifier. The data link connection identifier, however, is unambiguous only over a single virtual channel, and it may change in the node on transition to the next virtual channel. The bit 2 of the second address field octet, called a DE bit (Discard Eligibility Indicator), is also significant for the discarding of frames. In accordance with the CCITT recommendation, it is allowable to discard a frame e.g. under congestion conditions if the DE bit of the frame has been set to one. As the other bits in the FR frame are not relevant to the present invention, they will not be described more closely herein. References 2 and 4 mentioned above are referred to for a more detailed description.

At a subscriber node on the periphery of the network (Figure 2), the subscriber connections 14a, 14b, etc., (which in the example illustrated in Figure 2 are connected to the same node) are connected at first to an identification and control unit 61, which receives FR frames formed in bridges 13 (Figure 1). The identification and control unit 61 reads the data link connection identifier DLCI from the address field of the

frame, and forwards the frame to an input buffer 62₁...62_n corresponding to the virtual connection indicated by the identifier. Each data link has a specific selector S3 which selects frames from the input buffers of each virtual channel and forwards them further to a centralized router 16, which routes the frames further to a classification unit 43 of the correct data link (the figure shows only one outbound data link). The classification unit 43 reads the identifier DLCI from the address field of the frame and selects from table T the service class corresponding to the virtual channel indicated by the identifier. On the basis of the classification it has completed, the classification unit 43 applies each frame to an output buffer 64a, 64b or 64c corresponding to the service class of the frame. Each outbound data link thus has three output buffers, one for each service class. The selector S2 selects the frames from the service-class-specific output buffers 64a...64c and forwards them to the trunk connection.

Traffic transmitted by the subscribers over the FR network is thus buffered on the input side of the subscriber node specifically for each virtual connection. Incoming frames 39 are chained dynamically over each virtual connection. Depending on the service class of the virtual connection, the chain length has a predetermined allowable maximum value; this is smaller in service classes 1 and 3 and greater in service class 2. The selector S3 reads the buffers 62₁...62_n e.g. in proportion to the amount of traffic allocated to them, whereby the fairness principle is met.

When a node in the network is congested, the aim is to reduce traffic on the virtual connections running through it already at the subscriber node at the source end of the network; thus the frames will not load

the other resources of the network only to be discarded on reaching the congested node. According to the invention, this is carried out in such a way that the identification and control unit 61 of the subscriber node monitors the fill rate of each virtual-channel-specific buffer $62_1 \dots 62_n$ (in a manner known per se), and discards all the contents of a buffer if the buffer receives a frame while being full. At the same time, the identification and control unit discards this incoming frame (which may be, for instance, the first frame of a longer frame packet). The frames received thereafter are stored in the emptied buffer. In Figure 2, the operations of the identification and control unit for controlling and emptying the buffer are indicated with bidirectional arrows FC.

Discarding all the contents of the buffer at a time relieves congestion in the network considerably more than discarding only a few frames. The level of congestion in the network becomes considerably lower when the entire virtual-channel-specific buffer is emptied, the length of the buffer being typically dozens of frames. Emptying the buffer ensures that the virtual channel in question hardly loads the network at all for a while. This also makes it possible to avoid unnecessary loading of the network, as frames which would have to be discarded at a later stage because of congestion are not supplied to the network.

In view of the applications using the network, the method of the invention has the advantage that the frame loss probability is lower than before, as a result of which the load caused by retransmissions is also reduced. Discarding a whole buffer leads very close to a situation where an entire window of frames is discarded and must be retransmitted by the application protocol. It is therefore advantageous to adjust the

length of a virtual-channel-specific buffer to correspond substantially to the length of a window of frames sent at a time by an application using the network.

5 Figure 4 illustrates the processing of frames at a trunk node of the network. An FR frame 39 of the format described above is received at first by a classification unit 43 specific for each data link. The classification unit 43 reads the data link connection
10 identifier from the address field of the frame, and selects the service class corresponding to the virtual channel indicated by the identifier. The virtual channels and the respective service classes are stored in a table T. On the basis of the classification it has
15 completed, the classification unit 43 applies each frame to an input buffer 44a, 44b or 44c corresponding to the service class of the frame. Each inbound data link has thus three input buffers, one for each service class. Each data link has a specific selector S1 which selects
20 the frames from the service-class-specific buffers and forwards them within the node. On the output side of the trunk node, the frames are connected to an interface corresponding to their particular data link. At this interface they are supplied in accordance with the
25 service class selected on the input side of the node to one of the three service-class-specific output buffers 45a, 45b or 45c, from which the selector S2 reads the frames further to the trunk connection. Each outbound data link has thus three output buffers, one for each
30 service class. Alternatively, classification units may be provided separately for each data link even on the output side of the node, in which case classification data need not be transferred within the node.

35 Although the invention has been described above with reference to the examples shown in the accompanying

drawings, it will be obvious that the invention is not restricted to these examples, but can be modified within the scope of the inventive concept disclosed above and in the appended claims. In its details, the structure of a subscriber node, for example, may vary in many ways even though frames are discarded according to the principle of the invention. There may be, for instance, a separate fill rate control unit per each virtual-channel-specific buffer, said control unit transmitting data on the fill rate to a unit 61, which (in addition to identifying the virtual channel) merely empties the buffer.

Claims

1. A method for congestion management in an FR network comprising subscriber nodes to which subscribers (A...E) are connected over data links (14a...14e), said method comprising determining the virtual channel associated with a frame (39) to be transmitted when it is received from a subscriber connection at a subscriber node, c h a r a c t e r i z e d in that
- 5
- 10 - data is buffered at the input boundary of a subscriber node to virtual-channel-specific buffers ($62_1...62_n$),
- the fill rate of said buffers is monitored continuously, and
- 15 - when a frame is received in a virtual-channel-specific buffer which is full, substantially all the contents of the buffer are discarded.
2. A method according to claim 1, c h a r a c - t e r i z e d in that the length of the virtual-channel-specific buffer ($62_1...62_n$) is adjusted to correspond substantially to the length of the frame packet sent at a time by an application using the network.
- 20
3. A subscriber node of a frame relay network, to which node subscribers (A...E) of the network are connected over data links (14a...14e), said node comprising input buffers at the input boundary of the node, and output buffers at the output boundary of the node, and means (16) for relaying frames from a input buffer to a desired output buffer, c h a r a c t e r -
- 25
- 30 i z e d in that it comprises
- virtual-channel-specific buffers ($62_1...62_n$) at the input boundary,
- first means (61) for monitoring the fill rate
- 35 of said virtual-channel-specific buffers, and

- second means (61) responsive to said first means (61) and to the frame received for discarding the contents of a virtual-channel-specific buffer.

5 4. A subscriber node according to claim 1, characterized in that said first and second means are provided at the same identification and control unit (61), through which the frames received at the subscriber node are supplied to the virtual-channel-specific buffers ($62_1 \dots 62_n$).

1/2

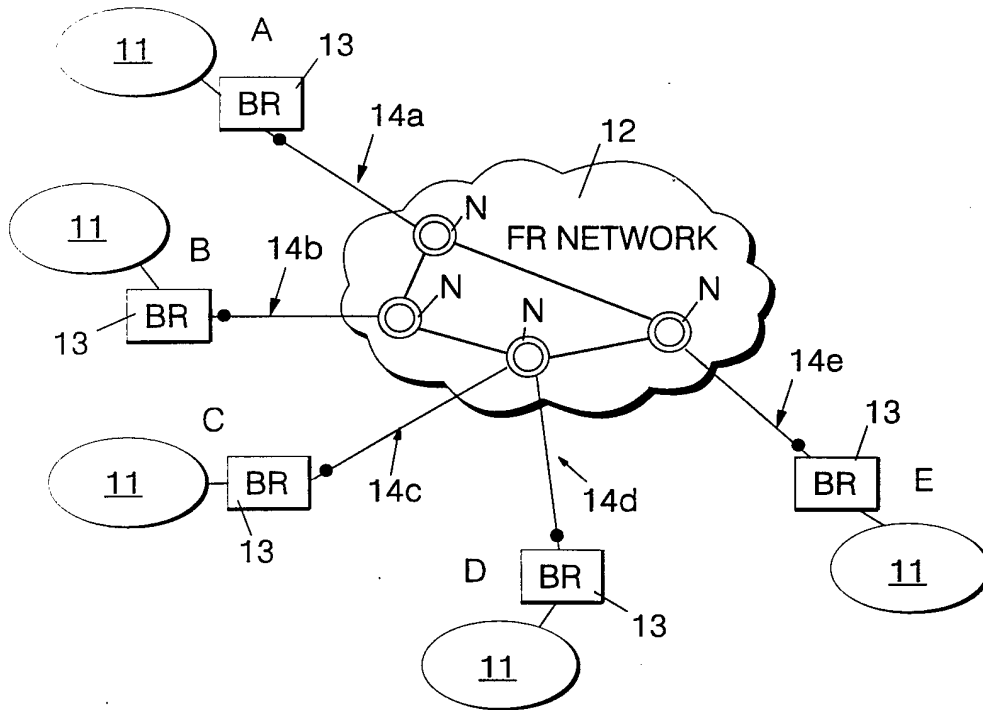


FIG. 1

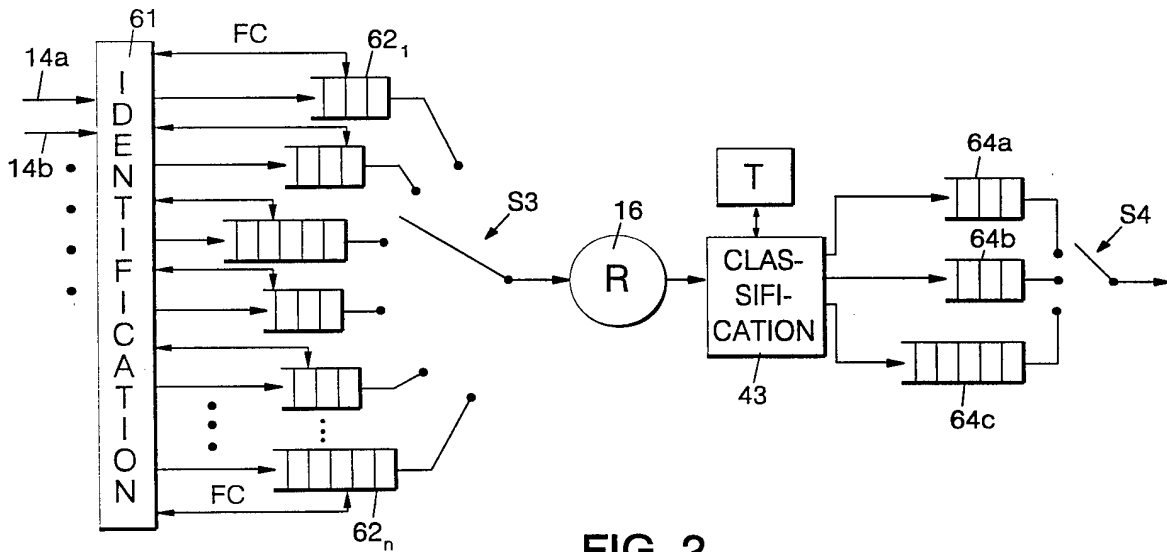


FIG. 2

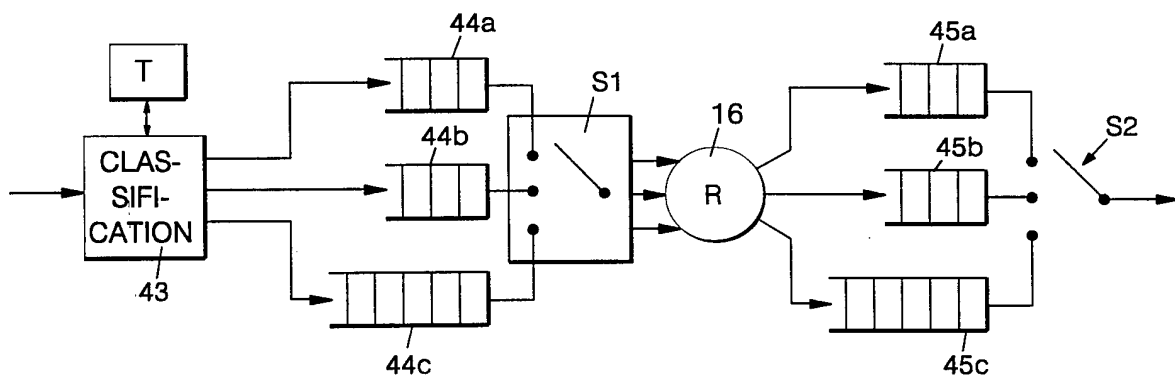
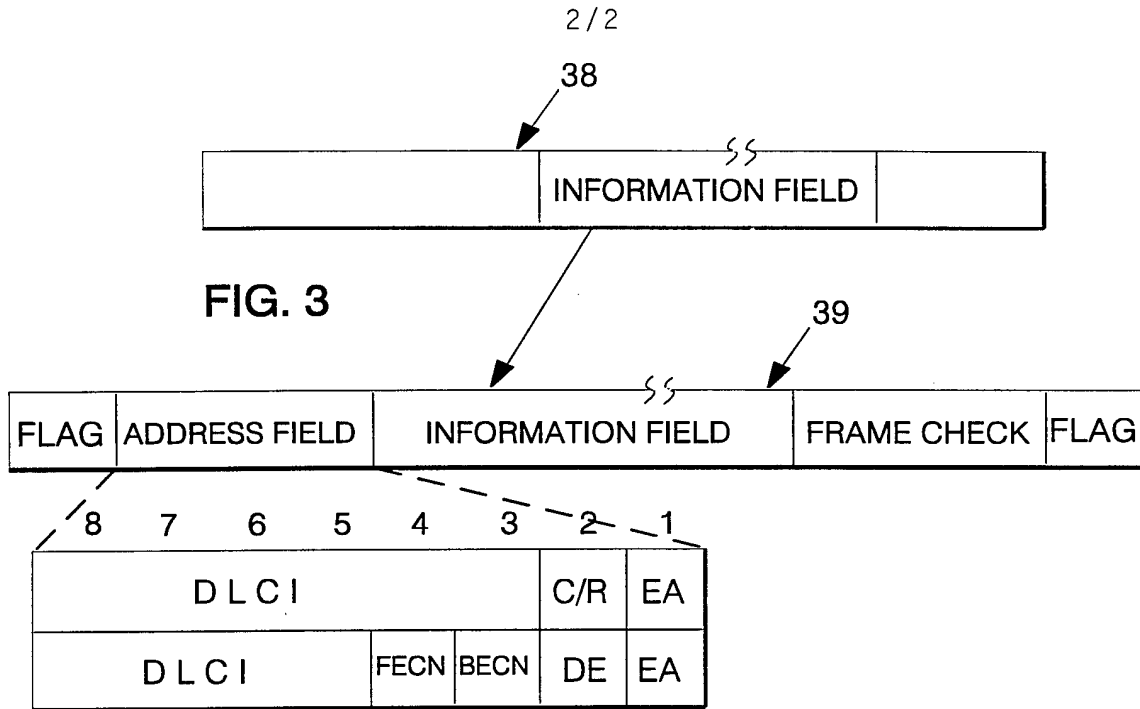


FIG. 4

INTERNATIONAL SEARCH REPORT

International application No.

PCT/FI 94/00534

A. CLASSIFICATION OF SUBJECT MATTER

IPC6: H04L 12/56

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC6: H04L

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

SE,DK,FI,NO classes as above

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

EPODOC, PAJ

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	EP, A1, 0499150 (HITACHI, LTD.), 19 August 1992 (19.08.92), column 2, line 43 - column 3, line 12; column 3, line 41 - line 50, abstract	1,3,4
Y	--	2
Y	US, A, 5163046 (ELLEN L. HAHNE ET AL), 10 November 1992 (10.11.92), abstract	2
A	--	1,3,4
A	WO, A1, 9314605 (CODEX CORPORATION), 22 July 1993 (22.07.93), page 3, line 10 - line 32	1-4
	-- -----	

☐ Further documents are listed in the continuation of Box C.☒ See patent family annex.

* Special categories of cited documents:

- "A" document defining the general state of the art which is not considered to be of particular relevance
- "E" earlier document but published on or after the international filing date
- "L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- "O" document referring to an oral disclosure, use, exhibition or other means
- "P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance: the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance: the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&" document member of the same patent family

Date of the actual completion of the international search

5 April 1995

Date of mailing of the international search report

06-04-1995

Name and mailing address of the ISA/
Swedish Patent Office
Box 5055, S-102 42 STOCKHOLM
Facsimile No. +46 8 666 02 86

Authorized officer

Adam Franzén
Telephone No. +46 8 782 25 00

INTERNATIONAL SEARCH REPORT
Information on patent family members

25/02/95

International application No.
PCT/FI 94/00534

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
EP-A1- 0499150	19/08/92	JP-A- 4257145 US-A- 5282203	11/09/92 25/01/94
US-A- 5163046	10/11/92	CA-A- 2029054 CA-A- 2030349 EP-A,A,A 0430570 EP-A,A,A 0430571 JP-A- 3186042 JP-A- 3188733 US-A- 5014265	31/05/91 31/05/91 05/06/91 05/06/91 14/08/91 16/08/91 07/05/91
WO-A1- 9314605	22/07/93	AU-B- 650195 AU-A- 3274193 EP-A- 0576647 JP-T- 6507290	09/06/94 03/08/93 05/01/94 11/08/94